Reasoning Patterns in Bayesian Games

An extended abstract of this paper was published in AAMAS 2011.

Dimitrios Antos
Harvard University
33 Oxford street 217
Cambridge, MA 02138
antos@fas.harvard.edu

Avi Pfeffer
Charles River Analytics
625 Mt. Auburn St.
Cambridge, MA 02138
apfeffer@cra.com

ABSTRACT

Game theory has been a useful in analyzing the structure and behavior of multi-agent systems. However, real world systems exhibit complexity and uncertainty which often impede the direct application of game-theoretic analyses. This often happens because agents lack common knowledge about the structure of the game or other agents' beliefs. Bayesian games have been traditionally employed to describe and analyze such situations. However, solving Bayesian games is computationally costly, and becomes even more so if the *common prior assumption* (CPA) has to be abandoned, which is sometimes necessary for a faithful representation of real-world systems. We propose using the theory of reasoning patterns in Bayesian games to circumvent some of these difficulties. The theory has been used successfully in common knowledge (non-Bayesian) games, both to reduce the computational cost of finding an equilibrium and to aid human decision-makers in complex decisions. In this paper, we first show that reasoning patterns exist for every decision of every Bayesian game, in which the acting agent has a reason to deliberate. This implies that reasoning patterns are a complete characterization of the types of reasons an agent might have for making a decision. Second, we show practical applications of reasoning patterns in Bayesian games, which allow us to answer questions that would otherwise not be easy in traditional analyses, or would be extremely costly. We thus show that the reasoning patterns can be a useful framework in analyzing complex social interactions.

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Miscellaneous

General Terms

Economics, Human Factors

Keywords

reasoning patterns, Bayesian games, game theory, Bayes-Nash equilibrium, heuristics

1. INTRODUCTION

Cite as: Reasoning Patterns in Bayesian Games, Dimitrios Antos, Avi Peffer, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Yolum, Tumer, Stone and Sonenberg (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. XXX–XXX.

The real world is a complex place, plagued with uncertainty. Designing agents to reason, make decisions and interact with others in such an environment is therefore a challenging problem. On the one hand, the number of states (or contingencies) that the agent needs to consider is prohibitively large. On the other hand, it often needs to interact with agents who have radically different beliefs about the situation unfolding. Some of them might have private information. Others might be employing fictitious or inaccurate beliefs about their opponents' and collaborators' utility functions or strategies. Finally, some agents might be boundedly rational or human, and therefore use heuristics, engage in limited reasoning or make simplifying assumptions to reach decisions faster. All the above problems are exacerbated when the agents need to make decisions quickly, and therefore cannot afford long-running computations. Also, agents need to be adaptive and perform well even if the situation changes, hence they cannot be employed with pre-computed optimal solutions fitting the particular narrow problem they face.

As an example, consider the game of poker. An agent designed to play poker with humans and computers must quantify and reason about its uncertainty over their utility functions (thus, for example, capturing their risk characteristics). It must be aware of the heuristic strategies (e.g., bluffing 15% of the time) that they might be using. It must understand that other players might be actively trying to elicit information about its own hand, by observing its betting behavior. And it must learn to classify its opponents, in order to understand who is novice and who might be more experienced. Finally, it needs to put itself in their shoes, and try to emulate their thought process and reasoning, e.g., predict the naive betting behavior of novices.

It becomes evident that our agents are required to reason in a quick, adaptive, yet accurate manner in the face of complexity and uncertainty. However, traditional game-theoretic approaches of dealing with the problem are often unsatisfactory. Situations in which players might disagree about the game being played are usually represented and analyzed as Bayesian games. In a Bayesian game, the private information of each player, as well as her beliefs about others' private information, are captured by her "type." Often, the common prior assumption (CPA) is also invoked, which requires the joint type (the vector of types of all the agents) to be drawn according to a probability distribution that is common knowledge. The CPA usually serves to simplify the game's representation and can be justified in some situations. However, the CPA is not always an appropriate modeling choice, especially in diverse populations of agents with different backgrounds in which agreement on a prior through repeated exposure is not warranted (see [11]) In a Bayesian game, agents are usually expected to adopt strategies comprising a Bayes-Nash equilibrium of the game.

This approach overlooks several issues. First, equilibrium solutions are hard to compute; there is no polynomial time algorithm for a Nash equilibrium in general games, and it is unlikely that one will be designed [8]. Second, a game usually has a multitude (or even an infinity) of equilibria, and there is no principled way to select one of them. Third, in Bayesian games without a common prior there are technical difficulties (e.g., infinite belief hierarchies) that make optimal solutions very expensive to compute. Also, equilibrium strategies might not be followed by human players, as experiments have demonstrated [7]. And finally, equilibria are mathematical solutions of an optimization problem, and hence leave the actual decision-maker "out of the loop." It is therefore hard to explain to a human why a particular strategy is good, or what exactly it accomplishes, beyond the dry fact that it "maximizes utility." In situations where the human feels responsible, needs to justify her decisions to others, or would like to have an intuitive understanding of her choices, prescribing equilibrium behavior might not suffice.

Reasoning patterns have been proposed as a framework that addresses these issues. In [12] the authors present four reasoning patterns, which are sets of features that capture the possible effects of an action on the acting agent's utility. The reasoning patterns can thus be used by the agent to conceptualize those effects, and thus be guided towards a "good" choice or strategy. In a sense, they capture what the agent needs to consider or remember when formulating her strategy. Experimentally, when humans are shown the reasoning patterns in a complex game, they make better decisions [3]. Moreover, if a decision of an agent does not have a reasoning pattern associated with it, we can safely ignore it for the purpose of computing a Nash equilibrium of the game [2], and thus reduce solving times. Theoretically, it has been shown that only four types of reasoning patterns exist and that these are "complete," in the sense that they can fully describe the relationship between action (decision) and effect (utility). We believe that reasoning patterns are thus useful in building good heuristics for decision-making in games.

However, the theory of reasoning patterns developed so far does not apply to Bayesian games. In this paper, we extend the theory to Bayesian games, distinguishing between games in which the CPA holds and those in which it does not. For the former, we show that the original definitions of the reasoning patterns are sufficient. For the latter, a new theory is required. This is because, for games in which there is no common prior, the different views of the agents cannot be combined in a single, unified model of the world, but instead reside in discrepant, inconsistent models. Hence, the effects of an action need to be traced across the agents' inconsistent models. Interestingly, we show that the reasoning patterns that might exist in such disconnected worlds are similar to the original four. Moreover, we prove that they are "complete," in the sense that any relation between an action and the agent's utility can be captured by some reasoning pattern. Finally, we show that these extended reasoning patterns can be used to capture interesting social interactions, and help answer questions that might otherwise be less obvious or very costly.

The structure of this paper is as follows. First, we review the reasoning patterns and past results. Next, we present Bayesian games, with and without a common prior. Following that, we present a structure to represent Bayesian games (of both types) that offers certain advantages in capturing the relationships among agents' beliefs. We then show how reasoning patterns can be identified in this structure, and present interesting classes of games for the analysis of social interactions. The paper concludes with a summary and extensions for future work.

Related Work

Our work aims at extending the ability for analyzing strategic situations beyond traditional game-theoretic analyses allow. Many avenues have been explored in that spirit. In [6] authors explore "cognitive hierarchies," a theory that suggests people engage in limited reasoning when analyzing a strategic situation. This can be used to circumvent computational issues with equilibrium calculation, although it usually assumes a distribution of the various hierarchy depths (steps of reasoning) people are expected to engage in. Team reasoning (see [13], [14]) seeks to replace individuals as the simplest reasoning unit with groups. The reasoning patterns, similarly, relate agents whose decisions influence one another. Finally, the field of epistemic game theory seeks to understand the relationship between rationality, players' belief in rationality, limited reasoning or knowledge, and game-theoretic reasoning and outcomes. The reasoning patterns aim at modeling reasoning at a coarser level than game-theoretic analyses, relaxing the assumptions made by traditional game-theory, yet circumventing the complexity or the paradoxes (e.g., see [5]) that rigorous epistemic game theory has revealed.

2. THE REASONING PATTERNS

When agents need to make a decision in a game, they need to be aware of the effects of their decisions. A good strategy is, essentially, a set of decisions (possibly stochastic) whose effects yield high total utility to the agent. These effects of decisions can be broken down into two categories: (a) effects on the deciding agent's utility function, and (b) effects on the other agents' knowledge, beliefs and actions. The former are direct ways of influencing one's payoff in the game, whereas the latter constitute indirect ways to achieve preferred outcomes.

Reasoning patterns capture the above two types of effects in a systematic way. The original paper [12] defines four reasoning patterns, and proves that these are "complete," in the sense that, if a decision of an agent cannot be associated with one of these four reasoning patterns, then the agent's choice of action bears no effect on her utility.

The four reasoning patterns are presented below. To give the reader an intuitive sense of these patterns, we present them within the context of the *poker* game.

1. Direct effect: This pattern captures the effect of a decision on that agent's utility *without* the intervention of any other agents. For instance, 'folding' during any round directly influences the payoff of a player in poker.

2. Manipulation: When agent i takes an action that will be (directly or noisily) observable by another agent j, then the latter might use this observation in his decision-making. We thus say that agent i has the ability to influence (manipulate) agent j's action through his own. For example, if j is using a heuristic strategy of the form "continue betting when others are betting, if you have three-of-a-kind or better," then agent i can influence j's decision by choosing to bet.

3. Signaling: Private information is often crucial in games. When parts of a decision-maker's knowledge is private to him (e.g., his utility function, personality type, or resources), then he might want to convey or conceal such information. Taking an action that will be visible to another agent j usually conveys something about the private information of agent i. Hence betting usually signals a good set of cards, or that one is bluffing.

4. Revealing-denying: As we mentioned before, information is crucial. Sometimes, an agent i will have the ability to control how much information will be accessible to another agent j by adopting a strategy. In poker, a player might engage in one more round of betting to explore whether his opponent is prepared to do so as well, thus causing his opponent to reveal something about his cards.

Technically, reasoning patterns (RPs) correspond to graphical properties of the Multi-Agent Influence Diagram (MAID) [10] representation of the game. In particular, each reasoning pattern is identified by a set of paths between nodes of certain types, satisfying certain d-separation properties. (A detailed technical definition of these paths can be found in the original paper.)

Notationally, we follow the standard graphical representation of MAIDs, whereby oval nodes represent probabilistic variables, rectangle nodes represent decisions belonging to one agent each, and diamond-shaped nodes represent utility functions (for example, see Figure 1). Arrows incoming to oval- or diamond-shaped nodes denote probabilistic dependencies (e.g., a utility node U is defined as a function of its parents in the graph). Arrows incoming to decision (rectangular) nodes represent "information." If decision node d belongs to agent i, then i is assumed to know the values of all variables that are parents of d in the graph. Finally, we make the common assumption of "perfect recall," whereby if decision node d' is a descendant of decision node d in the graph, and both of them belong to agent i, then i in d' is assumed to know the values of parents of both d and d', as well as the value for d that he chose in the past.

Due to the fact that reasoning patterns are "complete," decisions without a RP can safely be ignored for the purposes of computing a Nash equilibrium [2]. One can solve a reduced game (without these decisions), and then assign a fully-mixed uniform strategy to them, yielding an equilibrium for the original game. For games with certain independence relations, this can result in exponential savings in computation time.

Moreover, RPs have been used to aid human decision-makers. In [3], players in a repeated, multiple-round Bayesian game with private information who received a short description of the game's reasoning patterns were able to better understand the effects of their actions and thus perform better (score-wise), compared to players that did not have access to such "advice."

3. BAYESIAN GAMES

Bayesian games are used to analyze situations in which the game being played is not common knowledge. Instead, some players may hold a different payoff table or pure strategy set to be true. In a Bayesian game a player's beliefs include her knowledge of the game description (payoffs, strategies), as well as a probability distribution over the beliefs other agents might have. The set of beliefs held by player i is known as her type (t_i), and this is assumed to be drawn from a typeset T_i for that player, which can either be finite (discrete types) or infinite (a continuum of types).

What is often assumed in the analysis of Bayesian games is that the joint type (the vector comprising of every agent's type) is drawn according to a distribution that is common knowledge among the agents. This is known as the common prior assumption (CPA). The common prior assumption has several conceptual and computational advantages, yet cannot be indiscriminately applied to all games. One useful feature of the CPA is that the differences among players' beliefs can be explained away as differences in information. In particular, players are assumed to originally agree (on the common prior), but after each one receives some information in private (her type) their beliefs are now divergent (i.e., as posteriors derived by conditioning the common prior on the player's given type). Another useful feature of the CPA is that it makes inferences "consistent." Hence, if a piece of information is publicly announced, each agent will condition the same common prior in the same fashion. In other words, what agent i will think agent j of type t_j will believe after having observed e is exactly what j of that type will in fact believe.

However, invoking the CPA is inappropriate in many situations in which this original common belief is hard to justify. If agents have not been given this prior, but need to learn it by interacting in the environment, they may have had inadequate or limited exposure to it by the time our game-theoretic analysis begins. In that case, it is safer to assume that agents have differing priors. After all, the CPA is not a direct consequence of rationality (for a discussion of these issues see [11]). On the other hand, it has to be noted that the absence of a common prior significantly complicates the representation of games, because nested beliefs develop (often of an infinite hierarchy). Various formalisms, including I-POMDPs [9], have been developed to address the representational and computational difficulties arising out of a relaxation of the CPA.

4. REPRESENTING BAYESIAN GAMES

Recently a formalism has been developed for representing Bayesian games [4]. This representation is useful for our purposes, as it represents the views of players about the game being played as collections of MAIDs. This is helpful, since MAIDs capture the "story" behind a game and the probabilistic influence of a variable on others. We briefly review this formalism in this section.

The game is represented as a set B of blocks. A block $b \in B$ consists of a model $m(b)$ and beliefs $\beta(b)$. The model $m(b)$ is a fully-specified game in extensive form, or in MAID form. The beliefs $\beta(b)$ consist of $n(n-1)$ probability distributions, where n is the number of players in the game. The distribution p_{ij}^b over B refers to the belief of agent i in block b over the block agent j might be in. Hence $\beta(b)$ contains a distribution for each player i, for each other player $j \neq i$. We then define T_i, the typeset for player i, to be the set of blocks B, for every agent i. The interpretation is that each player i is assigned to a block in private, which represents her type; she knows her block (type), but other players do not. When player i is in block b, she holds the description of the game to be $m(b)$, and believes other players j to be assigned to blocks according to the $n-1$ distributions p_{ij}^b. Strategy sets, as well as Bayes-Nash equilibria, are defined in the same fashion as in every Bayesian game.

This formalism allows us to represent the dependencies among agents' beliefs graphically. The blocks in B are the nodes of this graph. Moreover, there is an edge from block b_1 to b_2, labeled Ôi.j,' if and only if $p_{ij}^{b_1}(b_2) > 0$, that is, if i in block b_1 believes that j might be using block b_2 with positive probability. We call this graph the belief graph. In [4] the authors show how the belief graph can be used to discover independencies and loosely-connected components among agents' beliefs.

As an example, suppose Bob lives in a place where the probability of rain every day is 20%. Bob prefers to carry an umbrella if rain is expected, but hates to carry it around on a sunny day. To make a decision, Bob consults the forecast on the morning paper, which is delivered to his place and is always very accurate in its predictions. However, Alice, his prank-loving daughter, who always wakes up before Bob, can replace the weather forecast section of the newspaper with a fake one of her liking, without Bob realizing the forgery. Alice gets tremendous pleasure from thinking of Bob getting wet in the rain. Suppose also that this is a one-shot game; after all, Alice cannot pull the same prank successfully every day.

To model this game, we shall have blocks b_1 and b_2 (see Figure 1). In the former, the weather 'forecast' is independent of the 'day's weather,' and depends solely on Alice's action. In the latter, the 'forecast' accurately depicts the 'day's weather,' and Alice's choice of action has no effect whatsoever. These two blocks represent two versions of the game. In both, Bob's utility depends on his action (whether he takes an umbrella), and the true weather. Also, in both blocks Alice's utility depends on Bob getting wet. Moreover, neither player can affect the true weather. Essentially, if a player is assigned to block b_1, he/she is aware of Alice's power to forge the newspaper section, whereas block b_2 captures a player who is oblivious to such trickery.

Blocks, however, also need to define the players' beliefs. We shall call the beliefs of i in block b "trivial" if, for all $j \in N - \{i\}$, it holds that $p_{ij}^b(b) = 1$. In our block b_2, then, it makes sense to assume that Bob's beliefs are trivial—Bob is oblivious to the possibility of Alice's switching of the forecast. In block b_1, similarly, Bob's beliefs ought to be trivial. As far as Alice is concerned, we can have her beliefs in b_2 be trivial. However, Alice may be uncertain over whether her father is truly oblivious of her plans, in which case it might be that $p_{AB}^{b_1}(b_2) = 0.7$. The belief graph for this game is shown in Figure 2.

What is interesting about this game is the relationship between Alice's and Bob's decisions, that emerges out of both blocks together, but is contained in neither separately. In block b_1, Bob knows that the forecast is not related to the

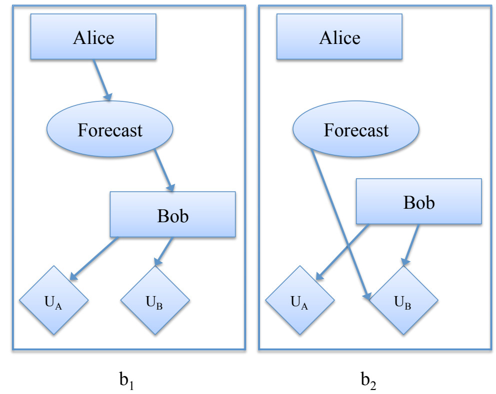

Figure 1: The MAIDs of our example game

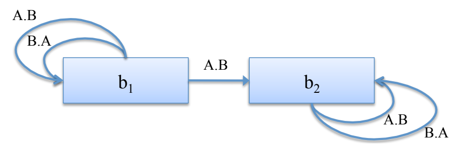

Figure 2: The belief graph of our example game

weather, so he will rationally ignore it; hence, Alice cannot affect Bob's action. In b_2, likewise, Bob will condition his action on the forecast, but Alice does not affect it, hence—again—there is no connection between the two players' actions. However, looking at both blocks together, we can see that Alice in b_1 thinks it is possible that Bob will condition on the forecast (if he uses b_2, which happens with probability 0.7), therefore Alice will want to affect it. Hence, for someone looking at the game "from the outside," Alice's action will have an effect on Bob's. How can we systematize these dependencies? The theory of reasoning patterns gives us a principled way to do this, as discussed in the next section.

5. REASONING PATTERNS ACROSS BLOCKS

As we have seen, in a Multi-Agent Influence Diagram (MAID), the directed edges represent probabilistic dependencies between variables (actions, events, utilities). We say that agent i "believes" in edge (x, y) if this edge exists in the block she has been assigned to. Clearly, if all the paths constituting a reasoning pattern exist wholly within the block of agent i, then that reasoning pattern is believed by i. However, a reasoning pattern may span several blocks.

First, we shall define a **motivated** decision. A decision d of agent i in block b is motivated if, when i has been assigned to block b, there is a strategy profile for all other agents σ_{-i} such that there are two actions a_1 and a_2 available in d that yield a different expected utility: $\mathrm{E}[u_i^b(a_1|\sigma_{-i})] \neq \mathrm{E}[u_i^b(a_2|\sigma_{-i})]$. In simpler terms, a decision is motivated for an agent if her action "matters" under some choice of strategy by other agents. (On the contrary, if a decision is not motivated, then any choice by i in that decision is equally good, so his choice "does not matter.")

Next, we define **well-distinguishing strategies**. A strategy σ_i for player i defines, for each decision d of i, a mapping from value assignments to its parents $\mathbf{Pa}(d)$ to a probability

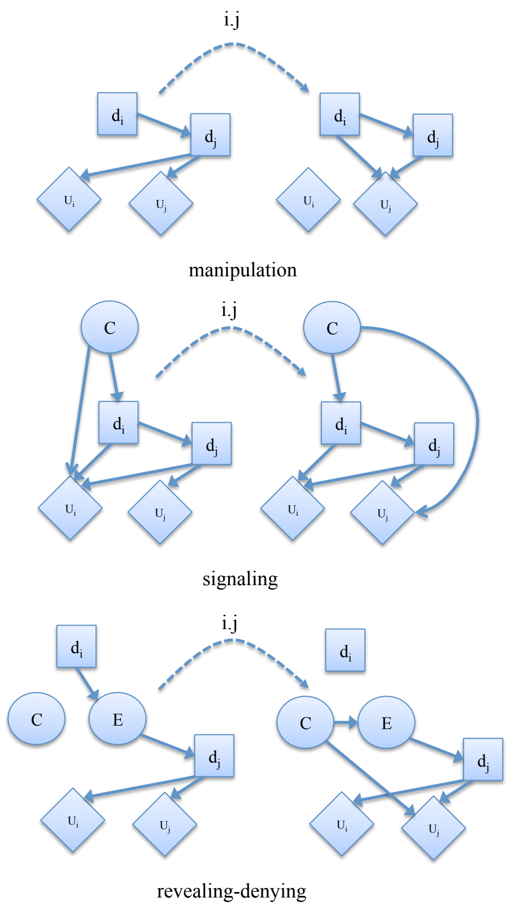

Figure 3: Examples of reasoning patterns across blocks

distribution over actions a_j available in d. A strategy σ_i is then well-distinguishing (WD) for i if, intuitively, i does not condition on a variable $v \in \mathbf{Pa}(d)$ if that variable has no effect on her utility (hence, "well-distinguishing")—for a formal definition see [12]. Hence, if configurations c_1 and c_2 of $\mathbf{Pa}(d)$ differ only in the value of variable v, a strategy would be well-distinguishing if it prevented i from choosing a different probability for any action a_j under c_1 and c_2, unless her expected utility was different, i.e., v "mattered."

We then define the four reasoning patterns, examples of which are given in Figure 3. In the definitions below we make the following assumption: For every node c in block b we assume there is a corresponding node c in every other block b', i.e., MAIDs in all blocks have the same set of nodes. Only edges are allowed to differ between blocks. In the following, we denote as $\mathbf{U}_i$ the set of i's utility nodes.

1. A decision d_i of agent i in block b has **direct effect** if there is a directed path wholly within $m(b)$ from d_i to one of the utility nodes of i that does not go through any other decision node (decision-free path). Direct effect captures the influence an agent's action has on her utility, without the intervention of other agents.

2. A decision d_i of agent i in block b has **manipulation** if there exists:

   (a) a directed, decision-free path in b from d_i to a decision d' of agent $j \neq i$,

   (b) a directed path from d' to a node in $\mathbf{U}_i$ in b that contains only motivated decision nodes (effective path), and

   (c) a directed effective path from d_i to a node in $\mathbf{U}_j$ (the utility nodes of j) that does not go through d' in some block b' for which $p^b_{ij}(b') > 0$, i.e., which is considered possible for j by i in b.

   Manipulation captures how an agent can influence her utility through exercising influence on another person's utility. Agent i takes an action that influences j's utility (path c). This changes j's optimization problem, because j stochastically knows i's action (path a). His (j's) optimal action under this setting can thus be influenced to increase i's utility (path b).

3. A decision d_i of agent i in block b has **signaling** if there exists:

   (a) a directed, decision-free path in b from d_i to a decision d' of agent $j \neq i$,

   (b) a directed effective[1] path from d' to a node in $\mathbf{U}_i$ in b,

   (c) an undirected effective path from a node C to a utility in $\mathbf{U}_j$ that is not blocked by the set $W^d_{d'} = \{d'\} \cup (\mathbf{Pa}(d') - \mathbf{Desc}(d))$, where $\mathbf{Desc}(d)$ are the descendants of node d, in some block b' for which $p^b_{ij}(b') > 0$,

   (d) a directed effective path from C to d_i in both b and b', and

   (e) an undirected effective path from C to a node in $\mathbf{U}_i$ in b, that is not blocked by W^C_d.

   Signaling captures the situation in which agent i can influence her utility through the action of another agent j, by conveying something about a variable in the world (C) that she can infer (path d) and which j cares about (path c).

4. A decision d_i of agent i in block b has **revealing-denying** if there exists:

   (a) a directed, decision-free path in b from d_i to a decision d' of agent $j \neq i$,

   (b) a directed effective path from d' to a node in $\mathbf{U}_i$ in b,

   (c) a directed effective path from a node E to d' in in some block b' for which $p^b_{ij}(b') > 0$,

   (d) a directed effective path from d_i to E in b,

   (e) a directed effective path from a node C to E in b and b', and

   (f) a directed effective path from C to a node in $\mathbf{U}_j$ in b' that is not blocked by $W^d_{d'}$.

   Revealing-denying captures the situation in which an agent (i) influences her utility thought the action of another agent (j) by controlling the uncertainty j has

[1] As explained above, a path is called "effective" if all the decision nodes on it are motivated.

over some variable (C) that he cares about. By increasing or decreasing the clarity by which j can infer C through E (paths d, e), agent j's optimal decision can be changed to benefit i's utility (path b).

We say that a decision node that has a reasoning pattern is "effective." We then prove the following theorem.

Theorem: *If a decision node d of agent i in block b is motivated, and all agents use WD strategies, then it is effective.*

PROOF: First, we look at our restriction to WD strategies, which implies the following for our graph structure. First, if a parent node v of a decision d has no effect on i's utility (for any assignment to the MAID paremeters), then is it d-separated from $\mathbf{U}_i$ given the set $W_d^v = \{d\} \cup (\mathbf{Pa}(d) - \{v\})$. Then, if i uses only WD strategies, she does not condition on such a variable v, therefore we can sever the edge (v, d) and retain an equivalent MAID graph.

Suppose then for the sake of contradiction that d is motivated but not effective. Since d is not effective, there is no directed, decision-free path from d to $\mathbf{U}_i$ in b (by definition of direct effect). Therefore, the only way d might affect $\mathbf{U}_i$ in b (in order to be motivated) is through some other agent $j \neq i$. Suppose d' is the decision node of agent j that facilitates this indirect effect. By our assumption, d must be able to affect d', therefore there must be a directed path from d to d' in block b; moreover, d' must be able to affect a node in $\mathbf{U}_i$, hence there must be a directed path from d' to $\mathbf{U}_i$ in block b.

But since j uses WD strategies, the path from d to d' can only exist if u, the parent of d' along that path, is not severed from d', that is, if u is not d-separated by $\mathbf{U}_j$ given $W_{d'}^u$. The only way this can be is if u is *d-connected* to $\mathbf{U}_j$ given this set, i.e., there is an undirected path from u to $\mathbf{U}_j$ that is not blocked by $W_{d'}^u$. This path should exist in the MAID of a block which is deemed likely for j to use, i.e., in a block satisfying $p_{ij}^b(b') > 0$.

Such a path might go through d or not. If it goes through d, there can be two cases: it either contains a sequence of nodes $< x, d, y >$ where y is a child of d, or y is a parent of d (by definition x must be a child of d). If y is a child of d, then the definition of manipulation holds, as there is now a subpath from d to $\mathbf{U}_j$ that does not go through d' by its blocking restrictions. If y is a parent of d, then signaling holds. This is because there is now a path from C, an ancestor of d (and y), to a node in $\mathbf{U}_j$. Moreover, since the edge from C to d is retained and not severed, there must be a path from C to some utility $\mathbf{U}_i$, which completes the definition for signaling. The only remaining case is if the path does not go through d at all. Then, there must be an edge E along that path that is a descendant of d, and a sequence of nodes along the path $< x, E, y >$, where y is a parent of E but not on the path from d to E. In that case, there must be a path from y to $\mathbf{U}_j$ that satisfies the blocking properties. But this is the definition for revealing-denying. As a final note, if any decision nodes exist along any of these undirected paths, the edges incoming to them must not be severed, therefore similar restrictions exist, i.e., the nodes must be effective. This is the reason why the definitions of the reasoning patterns use "effective paths." In conclusion, it cannot hold that a node is motivated but not effective, Q.E.D. $\square$

The theorem is significant because it renders the four reasoning patterns "complete." If a decision of an agent has no reasoning pattern, the theorem states that the decision will not be motivated, and thus the agent will have no reason to prefer one decision over another. Therefore, by examining the reasoning patterns we can capture all the reasons why an agent might choose an action.

Notice that if the CPA holds in a Bayesian game, we do not need to represent the game using multiple blocks. In this case, the joint type $\mathbf{T}$ can be a variable in a MAID, and individual types T_i can be deterministic children of $\mathbf{T}$. Then, we can draw an edge (T_i, d) for all decisions d of agent i, and our representation is complete. Since Bayesian games with a common prior are then representable in a single MAID, the original definition for the reasoning patterns can be used unchanged. However, as soon as the CPA is abandoned, it is no longer possible to use a single graph to represent the entire Bayesian game, and the new definitions need to be used instead. Of course, one can always represent a common-prior game in the block formalism, by making all agents' beliefs in every block the same (i.e., $p_{ij}^b \equiv p_{ij}^{b'}$ for all b, b' and all i, j). In that case, the new definitions yield the same set of reasoning patterns.

6. USING REASONING PATTERNS TO ANALYZE SOCIAL INTERACTIONS

We illustrate the usefulness of reasoning patterns in the analysis of Bayesian games by means of an example. Imagine there is an intelligence agency consisting of N agents. These agents collect information in the world, then summarize and interpret it, passing it on to their superiors, who then aggregate all the information and make decisions. Such a domain can be represented by a MAID. Rectangles are the actions taken by the agents, and oval nodes are information collected in the world or passed between agents. If all agents are cooperative, then all can be assumed to share a utility function U, which is represented as a single diamond node in the graph.

However, some of the agents might be "confederates." Such agents are trying to subvert the operation of the agency, and therefore can be assumed to have a different utility function U', which gives them a high value when the agency fails (i.e., when U is low). The agency is aware of the possibility of confederates among its members.

To take a simple case, suppose $N = 4$, named 1, 2, 3 and 4 respectively. Each agent i might be a confederate ($c(i) = 1$) or not ($c(i) = 0$). If agent i is a confederate, we also assume he knows all other agents that are confederates. Finally, we make the assumption that there are either zero or exactly two confederates in the agency.

In a Bayesian game, each agent would have a type t_i, drawn from a set T_i. Each type would have to indicate (a) whether the agent is a confederate, and (b) if the agent is indeed a confederate, the identity of the other confederate. Hence $T_i = \{(c, j) : c \in \{0, 1\}, j \in N \cup \{\emptyset\} - \{i\}\}$, with the restriction that, if $c = 0$ then $j = \emptyset$, and if $c = 1$ then $j \neq \emptyset$. The joint type vector $\mathbf{T} = \times_i T_i$ might be drawn from a common distribution $p(\mathbf{T})$, or not. The latter case, in which the common prior assumption does not hold, might be necessary to describe cases in which some agents trust the various members of the agency more than others, e.g., if the prior of 1 for the possibility of 3 being a confederate is different than the prior held by 2 for the same event.

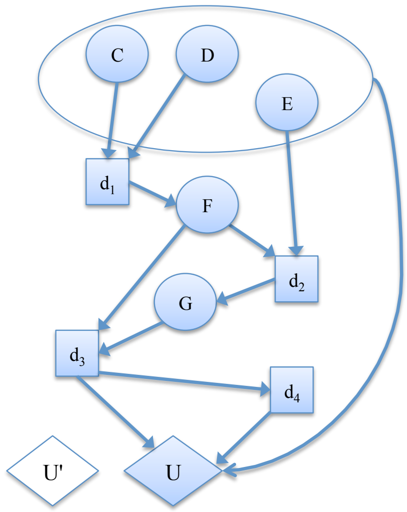

Figure 4: Agency example

Imagine now that the graph looks as in Figure 4. Agents 1 and 2 get information from the world (C, D and E) and compile reports (F and G). Agent 3 then makes a decision which is being communicated to agent 4, who makes a final decision. The utility node U is influenced by the decisions of agents 3 and 4, but is being shared by all agents. Agents in this game have reasoning patterns: Agent 4 has direct effect. Agent 3 has both direct effect (edge (d_3, U)) and manipulation (through 4). Agents 1 and 2 have signaling (as they signal the values of C, D and E) to agent 3.

Suppose now we were interested in answering the following question, set forth by agent 4, who is not a confederate: "Which pairs of agents should be more feared to be confederates?" and "Which pairs of agents are more likely to be the confederates, given that misreported information have been observed in node G?" In a traditional analysis, we would have to know, given $t_4 = (0, \emptyset)$, what the distribution of (t_1, t_2, t_3) is and, given this distribution, what the Bayes-Nash equilibria of the game are. Then, we would answer the first question by trying to compare the expected behavior of the players under the various Bayes-Nash equilibria with the observed behavior, as indicated by the misinformation received by player 4.

The problems with this analysis are that (a) there might be a multitude (or infinity) of equilibria, making the comparison hard, (b) equilibria are not easy to compute to begin with,[2] and (c) players might not agree on which equilibrium is to be played, or they might not be rational equilibrium players at all. Furthermore, this analysis requires that we know the probability distributions of all variables C, D, E, F and G, as well as the exact formula in U.

On the contrary, reasoning patterns allow us to do the following: First, we can represent this game in blocks $B = \{b_1 = (0, \emptyset), b_2 = (1, 2), b_3 = (1, 3), \ldots\}$. Each player i is assigned to one of the blocks $B - \{(1, i)\}$, because there must be exactly two confederates if he happens to be one himself. Edges between the blocks are drawn accordingly. Next, we can run the (polynomial) algorithm for the detection of reasoning patterns described in [1] and get a list of them.

We may then claim that the agents that have reasoning patterns such as manipulation, signaling and revealing-denying are more susceptible to being confederates than other agents. This requires some explaining. In signaling, the value being signaled must not be observable by other agents (otherwise the blocking conditions for path c in the definition do not hold). Hence, if an agent has a signaling reasoning pattern, that means he has the ability to misrepresent information without that being directly detectable. Similarly, through a revealing-denying reasoning pattern he controls access to information other agents have. Agents with this pattern can greatly enhance or impede the decision-making ability of their superiors. Likewise, manipulation reasoning patterns involve fabricating information that is input to some other agent's problem.

But the reasoning patterns do not just tell us that there might be an effect. They tell us "what the effect *is*," e.g., which variable is being signaled, or which variable will contain fabricated information. For instance, in the manipulation reasoning pattern, the confederate (i) will alter the value of the nodes on the path from d_i to d_j, where j is one of his superiors and no intermediates exist between i and j. Hence, if we receive evidence that one of these reports are fabricated, we can immediately cast suspicion upon agent i.

Also notice that the reasoning patterns analysis does not require knowledge of the exact utility function, or all the probabilistic dependencies. But if such knowledge is available, we may quantify the reasoning patterns, and calculate the expected utility of misrepresenting a variable by a confederate. Still, reasoning patterns would enable us to limit this search within the variables that the alleged confederate would have a reason to maliciously influence through his reasoning patterns.

7. CONCLUSION & FUTURE WORK

This paper addresses the issue of agent design for environments of high uncertainty and complexity, in which traditional game-theoretic solutions are inefficient or inadequate. We need agents to be quick and adaptive, and sometimes even explain and justify their decisions to humans. An alternative approach to equilibrium computation is offered by the theory of reasoning patterns, which was developed for non-Bayesian games in [12]. The reasoning patterns capture the possible reasons why an agent might take an action, by projecting the effects of her action on her utility. Our paper extends the definitions of the four reasoning patterns to Bayesian games, with and without a common prior. Under the common prior assumption (CPA), the fact that the game is representable in a single graph makes the original definition applicable. When the CPA is relaxed, however, new definitions are required. We prove that, just as in the common prior case, the reasoning patterns created by our definitions are complete, in that they capture all situations in which an agent is motivated to make a decision. We also show how reasoning patterns can be used to analyze interesting strategic interactions, even in situations where agents disagree about the definition of the game being played or focus on different aspects of it. For the future, we plan to (a) experimentally verify that reasoning patterns across blocks are practically useful for human decision-making, and (b)

[2]In this graph computation would be trivial, but in larger graphs it would be significantly more challenging.

develop a more precise, formal definition of particular social interactions that the reasoning patterns can capture.

8. REFERENCES

[1] D. Antos and A. Pfeffer. Identifying reasoning patterns in games. In *UAI*, 2008.

[2] D. Antos and A. Pfeffer. Simplifying games using reasoning patterns. In *AAAI*, 2008.

[3] D. Antos and A. Pfeffer. Using reasoning patterns to help humans make decisions in complex games. In *IJCAI*, 2009.

[4] D. Antos and A. Pfeffer. Representing bayesian games without a common prior. In *AAMAS*, 2010.

[5] A. Brandenburger. The power of paradox: some recent developments in interactive epistemology. *International Journal of Game Theory*, 2007.

[6] C. F. Camerer, T. H. Ho, and J. K. Chong. A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 2004.

[7] T. N. Cason and T. Sharma. Recommended play and correlated equilibria: a case study. *Economic Theory*, 2007.

[8] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a nash equilibrium. In *STOC*, 2006.

[9] P. Doshi and P. Gmytrasiewicz. On the difficulty of achieving equilibria in interactive pomdps. In *AAAI*, 2006.

[10] D. Koller and B. Milch. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, 2003.

[11] S. Morris. The common prior assumption in economic theory. *In Economics and Philosophy*, 1995.

[12] A. Pfeffer and Y. Gal. On the reasoning patterns of agents in games. In *AAAI*, 2007.

[13] R. Sugden. The logic of team reasoning. *Philosophical Explorations*, 2003.

[14] R. Sugden. Nash equilibrium, team reasoning and cognitive hierarchy theory. *Acta Psychologica*, 2007.